Efficient continuous-time quantum Monte Carlo algorithm for fermionic lattice models

Mauro Iazzi and Matthias Troyer
Theoretische Physik, ETH Zurich, 8093 Zurich, Switzerland

Efficient continuous time quantum Monte Carlo (CT-QMC) algorithms that do not suffer from time discretization errors have become the state-of-the-art for most discrete quantum models. They have not been widely used yet for fermionic quantum lattice models, such as the Hubbard model, nor other fermionic lattice systems due to a suboptimal scaling of $O(\beta^3)$ with inverse temperature β, compared to the linear scaling of discrete time algorithms. Here we present a CT-QMC algorithms for fermionic lattice systems that matches the scaling of discrete-time methods but is more efficient and free of time discretization errors. This provides an efficient simulation scheme that is free from the systematic errors opening an avenue to more precise studies of large systems at low and zero temperature.

PACS numbers:

arXiv:1411.0683v2 [cond-mat.str-el] 1 Jul 2015

Monte Carlo simulations of quantum systems are often performed using an imaginary time path integral formulation [1] to map the partition function of the quantum system to that of an equivalent classical one [2, 3]. These imaginary time paths, whose extent corresponds to the inverse temperature $\beta = 1/k_B T$ are then sampled using Monte Carlo methods. Path integrals are usually formulated on a discrete imaginary time mesh with nonzero time step Δ_τ in order to regularize the generally fractal paths. An extrapolation of the measured observables to $\Delta_\tau \to 0$ is then required to obtain accurate results corresponding to those of the original quantum system.

For discrete quantum lattice models, some quantum Monte Carlo (QMC) algorithms exist that are free from time discretization errors, such as Handscomb's method for Heisenberg spin models [4] or its generalization, the stochastic series expansion (SSE) algorithm [5]. They avoid an explicit introduction of time discretization by working with a Taylor expansion. For these and other discrete quantum lattice models one can also avoid time discretization errors in a path-integral formulation by realizing that the lattice structure already provides a regularization of the path integral.

Over the last two decades a new category of path-integral quantum Monte Carlo algorithms has thus been developed that work directly in the continuous time limit $\Delta_\tau \to 0$, removing the need for an extrapolation and often significantly speeding up the simulations. The first of these continuous-time quantum Monte Carlo (CT-QMC) algorithms have been for quantum spin systems and boson systems [6, 7]. Combined with efficient non-local update algorithms, such as cluster updates [8] or the worm algorithm [7] the gains in efficiency resulting from continuous time schemes are such that the simulation of unfrustrated spins and bosons is now considered a solved problem.

The generalization of CT-QMC to fermionic systems has been less straightforward, but has finally been achieved over the past decade by using time-dependent perturbation theory formulations of continuous time path integrals [9]. The first fermionic CT-QMC algorithm for lattice models [10] was followed by a number of algorithms for fermionic quantum impurity problems [9, 11–13]. These algorithms have been widely employed as quantum impurity solvers [9] *i.e.* for simulating an open system embedded in a non-interacting bath. They have replaced discrete time algorithms as the state of the art method by being significantly more efficient, avoiding the need to extrapolate in Δ_τ, and allowing the simulation of a much wider class of models. In particular, they have revolutionized the solution of the quantum impurity problem arising from self-consistent dynamical mean field (DMFT) theories [14–16] and their cluster extensions [17]. They allow the accurate simulation of much larger systems for Hubbard-type problems [18–20] and enable to go beyond density-density interactions by allowing the full Coulomb interaction to be included [21] thus opening the way to realistic materials simulations by multi-orbital DMFT [22].

Despite their enormous success for quantum impurity models, CT-QMC methods are rarely used for fermionic lattice models [18]. There, discrete time methods [23] are still the method of choice because of better scaling behavior. Existing fermionic CT-QMC algorithms all scale as $O(\beta^3 V^3)$ with the inverse temperature β and the lattice size V, since these algorithms require operations to be performed on square matrices with dimension $O(\beta V)$. For quantum impurity problems, which are described by time-dependent actions after integrating out the bath, also discrete time algorithms have the same scaling [24]. However, for quantum lattice models, discrete time algorithms exist that operate on β/Δ_τ matrices of dimension $O(V)$ and the effort thus scales only as $O(\beta V^3)$. The substantially reduced scaling ensured a significant competitive advantage of the discrete time approach.

Several (unpublished) attempts have been made to develop efficient CT-QMC methods for quantum lattice models. Naïve approaches have failed, giving either a worse sign problem or a $O(\beta V^4)$, erasing the advantage from the better scaling in temperature already for

medium size systems. In this Letter we show how to overcome the issues and present a CT-QMC algorithm that has linear scaling in β while retaining the cubic $O(V^3)$ complexity with respect to volume.

While our method is more general, we will – for the sake of simplicity – focus our presentation on the Hubbard model. The Hubbard model is the prototypical example of a strongly correlated fermionic system. It consists of spin-$\frac{1}{2}$ fermionic particles that can hop between neighbouring sites of a lattice and repel via a contact interaction. The full Hamiltonian H is given by a sum of the noninteracting and interacting parts $H = H_0 + H_I$

$$H_0 = -\sum_{x,y,\sigma} t_{xy} c^\dagger_{x\sigma} c_{y\sigma} \quad (1)$$

$$H_I = U \sum_x \hat{h}_x \equiv U \sum_x \left(\hat{n}_{x\uparrow} - \frac{1}{2}\right)\left(\hat{n}_{x\downarrow} - \frac{1}{2}\right). \quad (2)$$

Here the creation operators $c^\dagger_{x\sigma}$ introduces a new fermion at site x with spin σ and the annihilation operator $c_{x\sigma}$ likewise removes one particle, subject to the canonical anti-commutation relations $\{c^\dagger_{x\sigma}, c_{y\sigma'}\} = \delta_{xy}\delta_{\sigma\sigma'}$. The occupation number for site x is given by the density operator $\hat{n}_{x\sigma} = c^\dagger_{x\sigma} c_{x\sigma}$. The tunneling matrix $t_{xy} = t_{yx}$ [33] gives the energy associated with the hopping from site y to site x. When two particles of opposite spin are present in the same site, the they repel each other with energy U.

To perform the CT-QMC we perform a time-dependent perturbation expansion in H_I obtaining [9, 25]

$$Z = \operatorname{tr} e^{-\beta H} = \operatorname{tr}\left[e^{-\beta H_0} \mathcal{T} e^{-\int_0^\beta d\tau H_I(\tau)}\right] = \quad (3)$$

$$= \sum_{k=0}^{\infty} \frac{1}{k!} \operatorname{tr}\left[\mathcal{T} e^{-\beta H_0} \int_0^\beta d\tau_1 \ldots \int_0^\beta d\tau_k \prod_{i=1}^{k} (-H_I(\tau_i))\right]$$

where $H_I(\tau) = e^{\tau H_0} H_I e^{-\tau H_0}$ is the perturbation term H_I in the interaction representation. We then further expand the partition function as

$$Z = \sum_k \int_0^\beta d\tau_1 \int_{\tau_1}^\beta \ldots \int_{\tau_{k-1}}^\beta d\tau_k \sum_{x_1,\ldots,x_k} w(c) \quad (4)$$

where the factor $\frac{1}{k!}$ is taken care of by time ordering $\tau_1 < \tau_2 < \ldots < \tau_k$. $c = \{(x_1, \tau_1), \ldots, (x_k \tau_k)\}$ denotes a continuous time path integral configuration with k vertices and weight

$$w(c) = \operatorname{tr}\left[e^{-\beta H_0} \prod_{i=1}^{k} (-U \hat{h}_{x_k}(\tau_k))\right] \quad (5)$$

CT-QMC now proceeds by sampling from all possible configurations c according to their weight $w(c)$.

The structure of the factors in Eq. (5) allows the weight to be rewritten as [26]

$$w(c) = (-U)^k \det(1 + \mathbf{B}(c, \beta)), \quad (6)$$

where single particle propagator matrix $\mathbf{B}$ is given by

$$\mathbf{B}(c,\beta) = e^{-\beta \mathbf{H}_0} \prod_i \mathbf{h}(\tau_i, x_i) = \\ = e^{-(\beta-\tau_k)\mathbf{H}_0} \mathbf{h}(x_k) \ldots e^{-(\tau_2-\tau_1)\mathbf{H}_0} \mathbf{h}(x_1) e^{-\tau_1 \mathbf{H}_0} \quad (7)$$

where $\mathbf{H}_0$ is a $2V \times 2V$ matrix with elements $[H_0]_{x\sigma,y\sigma'} = t_{xy}\delta_{\sigma,\sigma'}$ and the matrix $\mathbf{h}(x_i)$ is given by

$$[\mathbf{h}(x_i)]_{x\sigma,y\sigma'} = \delta_{xx_i}\delta_{x_i y}\delta_{\sigma\uparrow}\delta_{\sigma'\uparrow} + \\ + \delta_{xx_i}\delta_{x_i y}\delta_{\sigma\downarrow}\delta_{\sigma'\downarrow} - \frac{1}{2}\delta_{xy}\delta_{\sigma\sigma'}, \quad (8)$$

and $\mathbf{h}(\tau_i, x_i) = e^{\tau_i \mathbf{H}_0} \mathbf{h}(x_i) e^{-\tau_i \mathbf{H}_0}$ is its time-displaced counterpart. In the case of the Hubbard the matrix $\mathbf{B}$ decomposes into two block matrices for each spin species, giving

$$w(c) = (-U)^k \det(1 + \mathbf{B}^\uparrow(c,\beta)) \det(1 + \mathbf{B}^\downarrow(c,\beta)) \quad (9)$$

The factor $(-U)^k$ introduces a sign problem for positive U, since any configuration with an odd number of vertices will have negative weight. On a bipartite lattice with nearest neighbour hoppings this trivial minus sign problem can be removed by mapping the repulsive model into an attractive one with interaction via a particle-hole transformation of the spin-down fermions $c_{x\downarrow} \to (-1)^x c^\dagger_{x\downarrow}$. This transformation changes the sign of U and thus removes this trivial sign problem, while the sublattice dependent sign avoids changing the sign of the kinetic energy. However a sign problems can still appear from the determinant.

This formulation of our algorithm is similar in spirit to the interaction-representation (CT-INT) algorithm for quantum impurity problems [9, 11] and can be considered a lattice CT-INT (or LCT-INT). Although not required, it has been found to be advantageous to instead use an auxilliary field decomposition [27] to remove the trivial sign for fermionic CT-QMC algorithms, leading the continuous time auxiliary field (CT-AUX) algorithm in the case of quantum impurity models [9, 13]. The CT expansion is applied with $\hat{h}(x) = (\hat{n}_{x\uparrow}\hat{n}_{x\downarrow} - 1)$. In our algorithm we can introduce an auxiliary field ρ giving an LCT-AUX representation

$$(1 - \hat{n}_{x\uparrow}\hat{n}_{x\downarrow}) = \frac{1}{2} \sum_{\rho=\pm 1} (1 + \rho\hat{n}_{x\uparrow})(1 - \rho\hat{n}_{x\downarrow}) \quad (10)$$

so that every vertex now also carries a spin degree of freedom and configurations are of the form $\{(x_i, \tau_i, \rho_i)\}_{i=1\ldots k}$. The auxiliary field dependent matrix $\mathbf{h}(x_i, \rho_i)$ now becomes

$$[\mathbf{h}(x_i, \rho_i)]_{x\sigma,y\sigma'} = \delta_{xy}\delta_{\sigma\sigma'} + \\ + \rho_i \delta_{xx_i}\delta_{x_i y}\delta_{\sigma\uparrow}\delta_{\sigma'\uparrow} - \rho_i \delta_{xx_i}\delta_{x_i y}\delta_{\sigma\downarrow}\delta_{\sigma'\downarrow}, \quad (11)$$

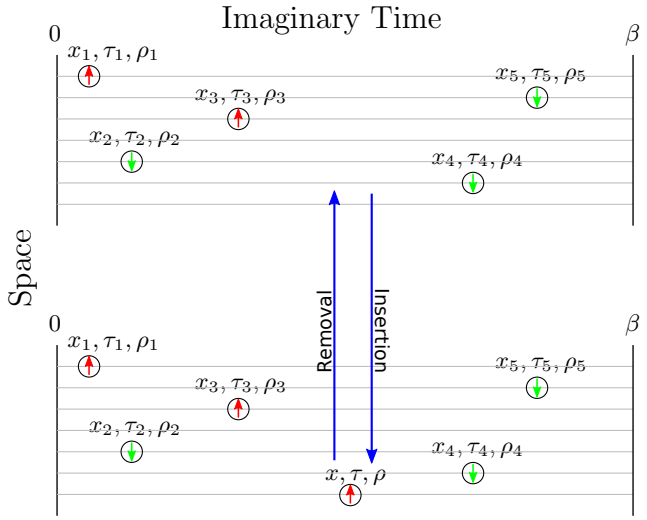

Figure 1: Sketch of the insertion and removal updates. The 5-vertex configuration c can be modified adding a vertex at site x and time τ with spin ρ. This leads to a proposed configuration c'. The reverse move consists of removing the vertex (x, τ, ρ). These two basic moves (insertion and removal) are sufficient to reach any term in the series (3) from any other.

and we end up with a weight similar to Eq. (6), but using $\mathbf{h}(x_i, \rho_i)$ instead of $\mathbf{h}(x_i)$.

To ergodically sample all possible terms in the series (3) it is sufficient to implement two types of Monte Carlo updates: insertion and removal of a vertex. They change the order k by ± 1 and are illustrated in Fig. 1. Starting from a configuration c with k vertices one proposes to randomly insert a new vertex with auxiliary field ρ at position x and time τ. The probability of accepting the new configuration c' is given, using the Metropolis algorithm [28], as $\min(1, R)$ with [34]

$$R = \frac{\beta V |U|}{k+1} \cdot \frac{\det\left[1 + \mathbf{B}(c', \beta)\right]}{\det\left[1 + \mathbf{B}(c, \beta)\right]}. \tag{12}$$

Conversely, the probability of removing a vertex is $\min(1, 1/R)$. The same acceptance ratio is derived for the LCT-INT version of the algorithm.

For any observable $\mathcal{O}$ one can write an estimator $O(c)$ which must be averaged to obtain an estimate of the quantum expectation value $\langle \mathcal{O} \rangle$. For equal-time observables such as densities, kinetic, and interaction energy, these are simple functions of the matrix $\mathbf{B}$. The single particle density matrix – or equivalently equal time Green function – $G(x, \sigma; y, \sigma') = \langle c_{x\sigma} c^\dagger_{y\sigma'} \rangle$ is estimated by measuring the matrix elements $\mathbf{G}_{x\sigma, y\sigma'}$ of the matrix

$$\mathbf{G} = \frac{\mathbf{B}(c, \beta)}{1 + \mathbf{B}(c, \beta)}. \tag{13}$$

The kinetic energy estimator is simply $E_0 = \mathrm{tr}(\mathbf{H}_0 \mathbf{G})$, while the interaction energy is given by $E_I = U \sum_x \mathbf{G}_{x\uparrow, x\uparrow} \mathbf{G}_{x\downarrow, x\downarrow}$.

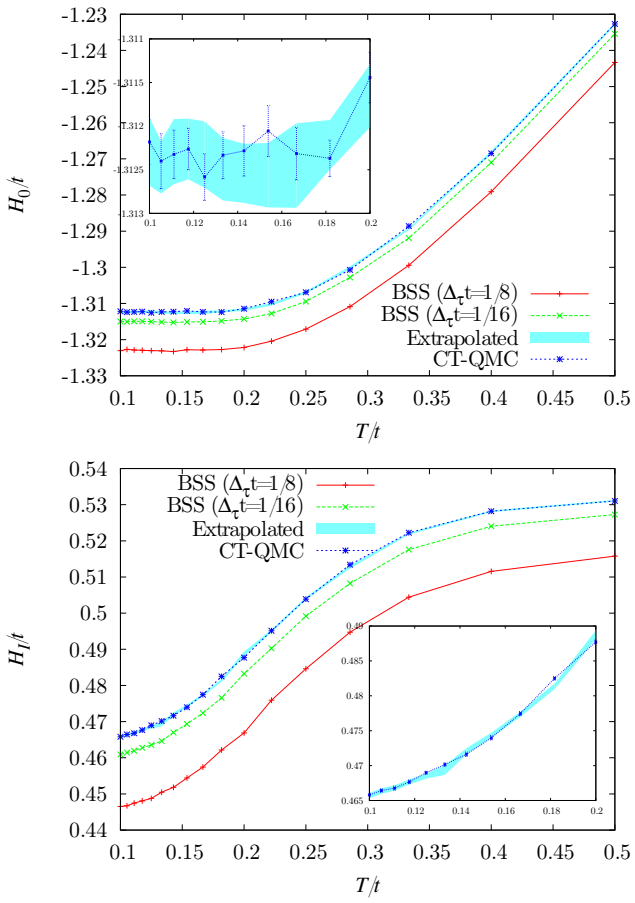

Figure 2: Kinetic energy and interaction energy for a 4×4 Hubbard plaquette at half filling with $U = 4t$, computed using the discrete time BSS algorithms for various values of $\Delta_\tau t$ and LCT-AUX. The inset compares the CT-QMC results to BSS results extrapolated to $\Delta_\tau \to 0$.

To demonstrate the reliability and performance of our algorithm we compared it to the discrete time BSS algorithm [23], which has so far been state of the art. Instead of starting from a continuous time representation (3), this algorithm is based on the Suzuki Trotter formula

$$e^{-\beta H} = \left(e^{-\Delta_\tau H_0} e^{-\Delta_\tau H_I}\right)^N + O(\beta \Delta_\tau^2). \tag{14}$$

which entails a so-called Trotter error due to time discretization that is quadratic in the time step $\Delta_\tau = \beta/N$. Figure 2 shows that the results from our CT-QMC algorithm agree perfectly with those obtained by extrapolating the finite-Δ_τ results obtained with the BSS algorithm to $\Delta_\tau \to 0$. The advantage of our algorithm is that it does not require this extrapolation in Δ_τ. This is particularly important for quantities such as the specific heat or double occupancy, which in the vicinity of phase transitions are very sensitive to the Trotter error. Equilibration and ergodicity features are also comparable, with autocorrelation times for the observables being very similar. This is easily understood as both the present method

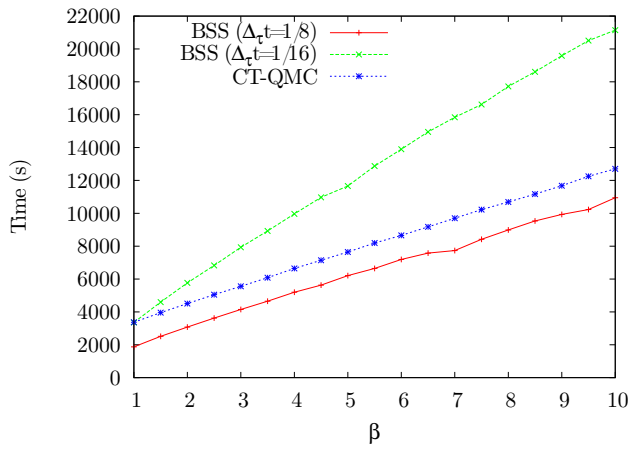

Figure 3: Computational time of the simulations from Fig. 2 as a function of β. Both the present method and discrete time algorithm have been run for 10^6 sweeps for thermalization then 10^6 sweeps while measuring observables. Both codes have been similarly optimized with delayed rank-1 updates. The linear dependence is clear in both pictures. The slope of the CT scheme is the same as the DT with $\Delta_\tau t = 1/8$ because the average number of vertices in a time slice of size t is around 120, i.e. roughly the same as $8V$=128. The CT code extrapolates to a much larger constant for $\beta \to 0$ mostly due to several allocations per update which have not been optimized.

and the state-of-the-art BSS scheme only employ local updates (vertex insertion/removal and single spin flip respectively).

The main computational effort in the algorithm is calculating the matrix $\mathbf{B}(\beta)$ and its changes when performing updates. Naïve multiplications of k matrices of dimension $O(V)$ would result in an effort of $kO(V^3)$. Considering that the number of vertices k grows with β, U and V we obtain a scaling of $O(\beta V^4)$, which is worse than the of the discrete time BSS algorithm. To achieve an overall $O(\beta V^3)$ scaling our algorithm works in the *eigenbasis* of H_0. Since the weight $w(c)$ is a determinant, it is unaffected by a basis change. Basis dependent quantities, such as observables, can be obtained via a rotation of the density matrix $\mathbf{G}$.

Diagonalizing $\mathbf{H}_0 = \mathbf{U}\mathbf{E}\mathbf{U}^\dagger$, where $\mathbf{E}$ is diagonal, the factors $e^{-(\tau_{i+1}-\tau_i)\mathbf{H}}$ also become diagonal matrices $e^{-(\tau_{i+1}-\tau_i)\mathbf{E}}$. The other factors are of the form $\delta_{kk'} + \gamma \mathbf{U}^\dagger_{kx_i}\mathbf{U}_{x_i k'}$, which is an identity matrix and an outer product of two vectors. Given this decomposition into sparse matrices and an outer product, the matrix multiplications can be performed with an effort $O(V^2)$, thus recovering the $O(\beta V^3)$ scaling of the BSS algorithm.

The product of matrices is in general an ill-conditioned matrix. To prevent numerical errors from creeping into the simulation, a stabilization procedure must be used, as explained in the Supplementary Material.

Common optimization techniques that are employed in other fermionic QMC algorithms can be applied here as well. Fast updates can be performed due to the fact that $\mathbf{B}(c)$ before the insertion (or removal) of a vertex, and $\mathbf{B}(c')$ after it, differ only by a single factor $\mathbf{h}(x,\tau)$, which is a diagonal matrix having all coefficients equal to 1 except a single one [23]. Several updates can also be combined in a delayed update scheme [19, 29]. We implemented delayed updates for both our discrete time and continuous time codes.

Our performance measurements in Fig. 3 confirm the linear scaling in β. We have compared LCT-AUX with a BSS code using the same set of optimizations. In these conditions LCT-AUX performs as well as the BSS algorithm with a commonly used time step of $\Delta_\tau = 1/8$. Since the BSS simulations have to be repeated for several values of Δ_τ and extrapolated, our unoptimized CT-QMC is already faster than a full discrete time calculation.

Using projections from a trial wave function our algorithm can be used for ground state simulations, similar to the discrete time algorithms [30] (See Supplementary Material). It can also be used as a a quantum impurity solver and used for DMFT calculations [14–17, 22, 31] by adding V_b non-interacting bath sites. The complexity of such an algorithm is $O(\beta U V (V+V_b)^2)$, which for low temperatures can be better than the $O(\beta^3 U^3 V^3)$ scaling of other CT-QMC algorithms [9, 11–13] for Hubbard-type models. The time-dependent Green functions $G(\tau; x, y) = \langle \mathcal{T}[c_x(\tau) c_y^\dagger(0)] \rangle$ required for DMFT can be measured using partial propagators $\mathbf{B}(c,\tau)$:

$$\mathbf{G}(\tau; \{x_i, \tau_i\}) = \frac{\mathbf{B}(c,\tau)}{1 + \mathbf{B}(c,\beta)} \tag{15}$$

with

$$\mathbf{B}(c,\tau) = e^{-\tau \mathbf{H}_0} \prod_{\tau_i < \tau} \mathbf{h}(\tau_i, x_i). \tag{16}$$

In general it is more stable to measure its Fourier transform, using non-uniform fast Fourier transformations [32].

The algorithm presented here is not specific to the Hubbard model. The only requirements are that the non-interacting Hamiltonian can be diagonalized once at the start of the algorithm to obtain the eigenvector matrix $\mathbf{U}$ and that the interacting Hamiltonian can be decomposed into exponentials of quadratic operators, i.e. $e^{\mathbf{A}_{xy} c_x^\dagger c_y}$ for some matrix $\mathbf{A}$. This is in general possible for any four-fermions interaction of the type $H_I = \sum_{kk'pp'} c_k^\dagger c_{k'} c_p^\dagger c_{p'}$ including the Coulomb interaction. As for the discrete time scheme, local interactions will retain the $O(V^3)$ scaling, but this might change in the general case (see Supplementary Material for more discussion).

In summary, we have presented a continuous time QMC algorithm for fermionic lattice models that has the same scaling as discrete time methods. This closes the

last prominent gap in the portfolio of CT-QMC algorithms, which have otherwise become the state of the art for bosons, quantum spins, and fermionic impurity problems. The main advantage of our algorithm is the absence of any time discretization error. This eliminates the need to either guess a small enough time step Δ_τ or extrapolate from multiple simulations at different Δ_τ to $\Delta_\tau \to 0$ and leads to shorter simulation times. Our algorithm can profit from the same numerical optimizations previously developed for other fermionic QMC algorithms [19, 23, 29] and can be used for finite temperature simulations, ground state calculations and quantum chemistry simulations and as a quantum impurity solver.

We acknowledge discussions with F.F. Assad and Lei Wang. This work was supported by the ERC Advanced Grant SIMCOFE, Microsoft Research, and the Swiss National Science Foundation through the National Competence Centers in Research NCCR QSIT and MARVEL. MT acknowledges hospitality of the Aspen Center for Physics, supported by NSF grant # 1066293.

Supplementary Material

Stabilization procedure

At the core of our algorithm is the calculation of the matrix

$$\mathbf{B}(c,\beta) = e^{-\beta \mathbf{H}_0} \prod_i \mathbf{h}(\tau_i, x_i) = \\ = e^{-(\beta-\tau_k)\mathbf{H}_0}\mathbf{h}(x_k)\dots e^{-(\tau_2-\tau_1)\mathbf{H}_0}\mathbf{h}(x_1)e^{-\tau_1\mathbf{H}_0} \quad (17)$$

As the ratio between the largest and lowest eigenvalue diverges information about the lowest eigenvalues and eigenstates is lost when the ratio between smallest and largest eigenvalues become of the order of the roundoff. Calculations of the determinant of $G(\beta)$ then becomes inaccurate.

Numerical stabilization of the product of matrices with an acceptable accuracy is made possible by periodically decomposing the intermediate result using a rank-revealing decomposition such as a singular value decomposition (SVD) or pivoting QR. One first multiplies all the vertices within a certain imaginary-time interval $\tau_S = \frac{\beta}{M}$, then decompose them

$$e^{-\Delta\tau \mathbf{H}_0} \prod_{\tau_i<\tau_S} \mathbf{h}(\tau_i, x_i) \to \mathbf{U}_1\mathbf{D}_1\mathbf{V}_1^t \quad (18)$$

The vertices up to $2\tau_S$ are then multiplied by $U_1 D_1$ and decomposed again.

$$\left(e^{-2\tau_S\mathbf{H}_0} \prod_{\tau_S<\tau_i<2\tau_S} \mathbf{h}(\tau_i, x_i)\right)\mathbf{U}_1\mathbf{D}_1 \to \mathbf{U}_2\mathbf{D}_2\mathbf{V}_2^t \quad (19)$$

The procedure is repeated until the full product has been performed

$$\mathbf{B}(c,\beta) = \mathbf{U}_M\mathbf{D}_M\mathbf{V}_M^t\dots\mathbf{V}_1^t. \quad (20)$$

The number of intervals M should be chosen so that the condition number of the partial products can be stored within machine precision. Additionally another SVD is performed whenever more than a certain number of vertices m is multiplied consecutively within the same interval.

An additional issue is present when the partial product $\mathbf{B}(c,\tau)$ has a high degeneracy (as is the case when few vertices are present). In this situation the numerical errors in the decompositions artificially lift the degeneracy and introduce spurious components in $\mathbf{B}$. To avoid this problem one can start the product from a random matrix $\mathbf{R}$ and finish it multiplying by $\mathbf{R}^{-1}$ obtaining $\mathbf{R}^{-1}\mathbf{B}\mathbf{R}$, which can be used in all the formulas for $\mathbf{B}$ with minimal modifications.

Fast updates

When a vertex is inserted at x', τ', a single matrix $\mathbf{h}$ is inserted in the expression for $\mathbf{B}$. This enables the use of rank-1 updates, via the matrix determinant lemma

$$\frac{\det(A+uv^t)}{\det(A)} = 1 + v^t(1+A)^{-1}u. \quad (21)$$

where A is a matrix and u, v two vectors. In this case the matrix is $\mathbf{B}$ and the two vectors are the x'-th column of the eigenvector matrix $\mathbf{U}$ evolved once backwards and once forward in time

$$u = e^{-(\beta-\tau_k)\mathbf{H}_0}\mathbf{h}(x_k)\dots e^{-(\tau_{j+1}-\tau_j)\mathbf{H}_0} \\ \mathbf{h}(x_j)e^{-(\tau_j-\tau')\mathbf{H}_0}\mathbf{U}_{kx'} \quad (22)$$

where τ_j is the first vertex after τ', and

$$v^t = \mathbf{U}^t_{x'k}e^{-(\tau'-\tau_{j-1})\mathbf{H}_0} \\ \mathbf{h}(x_{j-1})\dots e^{-(\tau_2-\tau_1)\mathbf{H}_0}\mathbf{h}(x_1)e^{-\tau_1\mathbf{H}_0} \quad (23)$$

The weight ratio can then be computed as

$$1 + v^t\mathbf{G}\mathbf{B}^{-1}u \quad (24)$$

where the vector $\mathbf{B}^{-1}u$ can be efficiently computed as

$$\mathbf{B}^{-1}u = e^{\tau_1\mathbf{H}_0}\mathbf{h}^{-1}(x_1)e^{(\tau_2-\tau_1)\mathbf{H}_0}\dots\mathbf{h}^{-1}(x_{j-1}) \\ e^{(\tau'-\tau_{j-1})\mathbf{H}_0}\mathbf{U}_{kx'} \quad (25)$$

if the newly inserted vertex is near the time origin. For this reason it can be advantageous to only allow vertex insertion (or removal) between the temporal origin and

a maximum time τ_w, and move the origin periodically, shifting all the vertex times by a constant time interval. The acceptance ratio becomes

$$R = \frac{\tau_w V|U|}{m+1}(1 + v^t \mathbf{G}\mathbf{B}^{-1}u) \tag{26}$$

where τ_w is the size of the time window where the updates can happen and m is the number of vertices within the same window. We can see that, if the vertex is added at the origin, the weight ratio reduces to the one for the discrete time case

$$\frac{w(c')}{w(c)} = -U\mathbf{G}_{x'x'} \tag{27}$$

This feature can be leveraged by always inserting or removing new vertices at the time origin, which is shifted at each update. The shifting procedure can be performed without a full recomputation of the matrix $\mathbf{G}$ using the formula

$$\mathbf{G}(\tau) = \mathbf{B}(\tau)\mathbf{G}(0)\mathbf{B}^{-1}(\tau) \tag{28}$$

which links the Green function at time τ to the Green function at time 0. If the number of vertices comprised between 0 and τ is m, the evolution matrices $\mathbf{B}(\tau)$ and its inverse can be applied with $O(mV^2)$ operations, so we need to limit the maximum time shift so that m is on average constant. Moreover, since this procedure is not numerically stable, one still needs to periodically recompute $\mathbf{B}$ and $\mathbf{G}$ from scratch.

Ground state method

As for the BSS algorithm, the present CT-QMC can be adapted to a *projector method* for ground state fermions. In such a scheme one wants to evaluate the expectation values

$$\langle O\rangle = \lim_{\theta\to\infty} \frac{\langle\psi_T|e^{-\frac{\theta}{2}\hat{H}}\hat{O}e^{-\frac{\theta}{2}\hat{H}}|\psi_T\rangle}{\langle\psi_T|e^{-\theta\hat{H}}|\psi_T\rangle} \tag{29}$$

where $|\psi_T\rangle$ is a trial wavefunction and $\lim_{\theta\to\infty} e^{-\theta\hat{H}}$ is the projector on the ground state. If the trial wavefunction is not orthogonal to the ground state Eq. (29) gives the correct results for the expectation value of O. The trial wavefunction is chosen to be a Slater determinant of p fermions described by the $p \times n$ matrix $\mathbf{P}$

$$|\psi_T\rangle = \prod_p (\mathbf{P}_{p1}c_1^\dagger + \ldots + \mathbf{P}_{pn}c_n^\dagger)|0\rangle \tag{30}$$

It is easy to recognize that θ has the same function as the inverse temperature β in the finite temperature scheme, so one can rewrite the expansion (3) as

$$\begin{aligned} Z &= \langle\psi_T|e^{-\theta H}|\psi_T\rangle = \langle\psi_T|e^{-\theta H_0}\mathcal{T}e^{-\int_0^\theta d\tau H_I(\tau)}|\psi_T\rangle = \\ &= \sum_{k=0}^{\infty}\frac{(-1)^k}{k!}\langle\psi_T|e^{-\theta H_0}\mathcal{T}\int_0^\theta d\tau_1\ldots\int_0^\theta d\tau_k\prod_{i=1}^k H_I(\tau_i)|\psi_T\rangle = \\ &= \sum_k\int_0^\theta d\tau_1\int_{\tau_1}^\theta\ldots\int_{\tau_{k-1}}^\theta d\tau_k\sum_{x_1,\ldots,x_k} w(c) \end{aligned} \tag{31}$$

where the weight is now given only by the expectation value of the operator $B(c,\theta)$ over the trial state $|\psi_T\rangle$. Since the trial state is a Slater determinant, the weight can still be expressed as a determinant of a single particle matrix

$$w(c) = \det(\mathbf{P}\mathbf{B}(c,\theta)\mathbf{P}^\dagger). \tag{32}$$

This weight can be computed using the same techniques s in the finite temperature case, however there are additional simplifications that can be exploited in the ground state formalism. If one splits the evolution in two parts, from 0 to τ and from τ to θ, we can define

$$\mathbf{L}(\tau) = \mathbf{P}\mathbf{B}(c;\theta,\tau) \tag{33}$$
$$\mathbf{R}(\tau) = \mathbf{B}(c;\tau,0)\mathbf{P}^\dagger \tag{34}$$

Performing a SVD of both $\mathbf{L}$ and $\mathbf{R}$ yelds

$$\mathbf{L} = \mathbf{V_L}\mathbf{D_L}{\mathbf{U_L}}^\dagger \tag{35}$$
$$\mathbf{R} = \mathbf{U_R}\mathbf{D_R}{\mathbf{V_R}}^\dagger \tag{36}$$

where $\mathbf{V_{L/R}}$ are unitary $p\times p$ matrices, $\mathbf{D_{L/R}}$ are diagonal with p elements and $\mathbf{U_{L/R}}$ are $n\times p$ matrices whose columns are orthonormal vectors. The equal-time Green function at time τ is then given by

$$\begin{aligned} \mathbf{G}(\tau) &= 1 - \mathbf{R}(\tau)[\mathbf{L}(\tau)\mathbf{R}(\tau)]^{-1}\mathbf{L}(\tau) = \\ &= 1 - \mathbf{U_R}[{\mathbf{U_L}}^\dagger\mathbf{U_R}]^{-1}{\mathbf{U_L}}^\dagger \end{aligned} \tag{37}$$

Since one needs only $\mathbf{U_{L/R}}$ to compute the Green function, the stabilization procedure is considerably simplified, since one can discard the intermediate $\mathbf{D}$s and $\mathbf{V}$s. The fast-update method can still be implemented in this formalism by noting that

$$\frac{\det(\mathbf{L}(\tau)(1+uv^t)\mathbf{R}(\tau))}{\det(\mathbf{L}(\tau)\mathbf{R}(\tau))} = 1 + v^t(1-\mathbf{G}(\tau))u \tag{38}$$

which again only requires the knowledge of the $\mathbf{U}$s to be computed.

Implicit and explicit bath

For Dynamical Mean Field Theory (DMFT) and Dynamical Cluster Approximation (DCA) calculations, it is

necesary to simulate a system where the noninteracting particle Green function includes a self energy $\mathbf{\Sigma}$

$$\mathbf{G}(i\omega) = \frac{1}{i\omega - \mathbf{H}_0 + \mathbf{\Sigma}(i\omega)} \tag{39}$$

where the self energy is calculated from previous iterations of the algorithm (see Ref. [] for a review of DMFT/DCA). This is the so-called *implicit bath* scheme, as the effect of a thermal bath is simulated indirectly. Since the self-energy depends in general on the imaginary frequency, the noninteracting systems is not described by a Hamiltonian and thus the present method cannot be employed. Instead one can use CT-AUX and CT-INT or the Hisrch-Fye QMC (which is of discrete-time type), which have all the same scaling properties.

Another option to simulate a system with a given self-energy is to introduce a number N_b of ancillary sites to the model, with creation and annihilation operators $a^\dagger_{y\sigma}$ and $a_{y\sigma}$. The dynamics of these bath sites and their interaction with the target system are described by the auxiliary Hamiltonian

$$\hat{H}_b = \sum_{y}^{N_b}\sum_{\sigma} \epsilon_{y\sigma} a^\dagger_{y\sigma} a_{y\sigma} + \sum_{x\sigma,y\sigma'} V_{x\sigma,y\sigma'}(c^\dagger_{x\sigma} a_{y\sigma'} + a^\dagger_{y\sigma'} c_{x\sigma}) \tag{40}$$

The values of the bath energy levels ϵ_y and hopping parameters V_{xy} are chosen so that the self energy induced by the bath into the system

$$\mathbf{\Sigma}_b(i\omega) = \sum_{xx'y} \frac{V_{xy}V_{x'y}}{i\omega - \epsilon_y} \tag{41}$$

fits the desired self-energy $\mathbf{\Sigma}$. The fit can be performed minimizing the Green function misfit

$$\chi^2(\{\epsilon, V\}) = \sum_{i\omega} w(i\omega) \left| \frac{1}{i\omega - \mathbf{H}_0 - \mathbf{\Sigma}(i\omega)} + \right. \\ \left. - \frac{1}{i\omega - \mathbf{H}_0 - \mathbf{\Sigma}_b(i\omega)} \right|^2 \tag{42}$$

with respect to the bath parameters. The weighting factors $w(i\omega)$ can be used to prioritize accuracy of the fit over relevant (usually lower) frequencies. Such *explicit bath* scheme is indeed applicable to the present method.

Local and non-local interactions

The present method can be applied to both local and non-local interactions. For a generic density-density interaction of the form

$$H_I = \sum_x \sum_y U(|x-y|) \left(\hat{n}_x - \frac{1}{2}\right)\left(\hat{n}_y - \frac{1}{2}\right) \tag{43}$$

we obtain several types of vertices coresponding to all possible pairs of sites (x, y) and the corresponding single-particle matrix $\mathbf{h}(x_i, y_i)$ will depend on both site indices. In contrast to the simple Hubbard model, each vertex has a different coefficient $U(|x_i - y_i|)$. Moreover in this case there will be vertices corresponding to interactions between same-spin particles. In such cases the matrix $\mathbf{h}(x_i, y_i)$ only acts on one component of the spin (e.g. only the up component) as a *rank*-2 matrix of the form

$$[\mathbf{h}(x_i, y_i)]_{x\sigma,y\sigma'} = \delta_{xx_i}\delta_{x_i y}\delta_{\sigma\uparrow}\delta_{\sigma'\uparrow} + \\ + \delta_{xy_i}\delta_{y_i y}\delta_{\sigma\uparrow}\delta_{\sigma'\uparrow} - \frac{1}{2}\delta_{xy}\delta_{\sigma\sigma'}, \tag{44}$$

i.e. a diagonal matrix with all coefficients equal to $-\frac{1}{2}$ except for the ones at x_i and y_i which have the sign flipped. In the energy eigenbasis this matrix becomes

$$[\mathbf{h}(x_i, y_i)]_{p\sigma,q\sigma'} = \mathbf{U}^\dagger_{px_i}\mathbf{U}_{x_i q}\delta_{\sigma\uparrow}\delta_{\sigma'\uparrow} + \\ + \mathbf{U}^\dagger_{py_i}\mathbf{U}_{y_i q}\delta_{\sigma\uparrow}\delta_{\sigma'\uparrow} - \frac{1}{2}\delta_{pq}\delta_{\sigma\sigma'} \tag{45}$$

The acceptance ratio for insertion of such a vertex becomes

$$R = \frac{\beta N_p |U|}{k+1} \cdot \frac{\det\left[1 + \mathbf{B}(c', \beta)\right]}{\det\left[1 + \mathbf{B}(c, \beta)\right]}. \tag{46}$$

where N_p is the number of possible pairs that can be inserted. In the case of nearest neighbour interactions on a square lattice $N_p = 8V$, with $2V$ coming all possible links and a factor 4 for all possible spin combinations. For a generic all-to-all interaction $N_p = 4V * (V-1) + V$.

The number of vertices in the simulation is proportional to the ground state interaction energy $\langle H_I \rangle$. As such it will still grow linearly for any local and quasi-local interaction, as well as with power law interactions decaying faster than $\frac{1}{V}$, resulting again in a scaling of $O(\beta U V^3)$ for the simulation. For Coulomb potentials the interaction energy can grow faster than V i.e. be super-extensive and results in a worse scaling with volume.

[1] R. P. Feynman, "Atomic Theory of the λ Transition in Helium," *Phys. Rev.*, vol. 91, pp. 1291–1301, Sept. 1953.

[2] M. Suzuki, S. Miyashita, and A. Kuroda, "Monte carlo simulation of quantum spin systems. i," *Progress of Theoretical Physics*, vol. 58, no. 5, pp. 1377–1387, 1977.

[3] J. A. Barker, "A quantum-statistical Monte Carlo method; path integrals with boundary conditions," *The Journal of Chemical Physics*, vol. 70, no. 6, p. 2914, 1979.

[4] D. C. Handscomb, "The Monte Carlo method in quantum statistical mechanics," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 58, pp. 594–598, Oct. 1962.

[5] A. W. Sandvik and J. Kurkijärvi, "Quantum Monte-Carlo Simulation Method for Spin Systems," *Physical Review B*, vol. 43, no. 7, pp. 5950–5961, 1991.

[6] B. B. Beard and U.-J. Wiese, "Simulations of discrete quantum systems in continuous euclidean time," *Phys. Rev. Lett.*, vol. 77, pp. 5130–5133, Dec 1996.

[7] N. Prokof 'ev, B. V. Svistunov, and I. S. Tupitsyn, "Exact, complete, and universal continuous-time worldline Monte Carlo approach to the statistics of discrete quantum systems," *Journal of Experimental and Theoretical Physics*, vol. 87, pp. 310–321, Aug. 1998.

[8] H. Evertz, "The loop algorithm," *Advances in Physics*, vol. 52, pp. 1–66, Jan. 2003.

[9] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, "Continuous-time monte carlo methods for quantum impurity models," *Rev. Mod. Phys.*, vol. 83, pp. 349–404, May 2011.

[10] S. Rombouts, K. Heyde, and N. Jachowicz, "Quantum monte carlo method for fermions, free of discretization errors," *Physical Review Letters*, vol. 82, no. 21, pp. 4155–4159, 1999.

[11] A. N. Rubtsov, V. V. Savkin, and A. I. Lichtenstein, "Continuous-time quantum Monte Carlo method for fermions," *Physical Review B*, vol. 72, p. 035122, July 2005.

[12] P. Werner, A. Comanac, L. de' Medici, M. Troyer, and A. J. Millis, "Continuous-Time Solver for Quantum Impurity Models," *Physical Review Letters*, vol. 97, p. 076405, Aug. 2006.

[13] E. Gull, P. Werner, O. Parcollet, and M. Troyer, "Continuous-time auxiliary-field Monte Carlo for quantum impurity models," *Europhysics Letters*, vol. 82, p. 57003, May 2008.

[14] W. Metzner and D. Vollhardt, "Correlated Lattice Fermions in d=∞ Dimensions," *Physical Review Letters*, vol. 62, pp. 324–327, Jan. 1989.

[15] A. Georges and G. Kotliar, "Hubbard model in infinite dimensions," *Physical Review B*, vol. 45, pp. 6479–6483, Mar. 1992.

[16] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, "Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions," *Reviews of Modern Physics*, vol. 68, pp. 13–125, Jan. 1996.

[17] T. Maier, M. Jarrell, and M. Hettler, "Quantum cluster theories," *Reviews of Modern Physics*, vol. 77, pp. 1027–1080, Oct. 2005.

[18] S. Fuchs, E. Gull, L. Pollet, E. Burovski, E. Kozik, T. Pruschke, and M. Troyer, "Thermodynamics of the 3d hubbard model on approaching the néel transition," *Phys. Rev. Lett.*, vol. 106, p. 030401, Jan 2011.

[19] E. Gull, P. Staar, S. Fuchs, P. Nukala, M. S. Summers, T. Pruschke, T. C. Schulthess, and T. Maier, "Submatrix updates for the continuous-time auxiliary-field algorithm," *Phys. Rev. B*, vol. 83, p. 075122, Feb 2011.

[20] J. P. F. LeBlanc and E. Gull, "Equation of state of the fermionic two-dimensional hubbard model," *Phys. Rev. B*, vol. 88, p. 155108, Oct 2013.

[21] P. Werner and A. J. Millis, "Hybridization expansion impurity solver: General formulation and application to Kondo lattice and two-orbital models," *Physical Review B*, vol. 74, p. 155107, Oct. 2006.

[22] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, "Electronic structure calculations with dynamical mean-field theory," *Rev. Mod. Phys.*, vol. 78, pp. 865–951, Aug 2006.

[23] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, "Monte carlo calculations of coupled boson-fermion systems. i," *Phys. Rev. D*, vol. 24, pp. 2278–2286, Oct 1981.

[24] J. E. Hirsch and R. M. Fye, "Monte carlo method for magnetic impurities in metals," *Phys. Rev. Lett.*, vol. 56, pp. 2521–2524, Jun 1986.

[25] N. Prokof'ev, B. Svistunov, and I. Tupitsyn, "Exact quantum monte carlo process for the statistics of discrete systems," *Journal of Experimental and Theoretical Physics Letters*, vol. 64, no. 12, pp. 911–916, 1996.

[26] S. Koonin, D. Dean, and K. Langanke, "Shell model monte carlo methods," *Physics Reports*, vol. 278, no. 1, pp. 1 – 77, 1997.

[27] J. E. Hirsch, "Discrete hubbard-stratonovich transformation for fermion lattice models," *Phys. Rev. B*, vol. 28, pp. 4059–4061, Oct 1983.

[28] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines," *The Journal of Chemical Physics*, vol. 21, no. 6, pp. 1087–1092, 1953.

[29] G. Alvarez, M. Summers, D. Maxwell, M. Eisenbach, J. Meredith, J. Larkin, J. Levesque, T. Maier, P. Kent, E. D'Azevedo, and T. Schulthess, "New algorithm to enable 400+ TFlop/s sustained performance in simulations of disorder effects in high-T c superconductors," in *2008 SC - International Conference for High Performance Computing, Networking, Storage and Analysis*, p. 10 pp., 2008 2008. 2008 SC - International Conference for High Performance Computing, Networking, Storage and Analysis, 15-21 Nov. 2008, Austin, TX, USA.

[30] G. Sugiyama and S. Koonin, "Auxiliary field monte-carlo for quantum many-body ground states," *Annals of Physics*, vol. 168, no. 1, pp. 1 – 26, 1986.

[31] E. Khatami, C. R. Lee, Z. J. Bai, R. T. Scalettar, and M. Jarrell, "Cluster solver for dynamical mean-field theory with linear scaling in inverse temperature," *Phys. Rev. E*, vol. 81, p. 056703, May 2010.

[32] L. Greengard and J. Lee, "Accelerating the nonuniform fast fourier transform," *SIAM Review*, vol. 46, no. 3, pp. 443–454, 2004.

[33] Or more generally $t_{xy} = t^*_{yx}$ in the presence of a gauge field. Since gauge fields generally introduce a phase problem we restrict ourselves to the case of real t_{xy}

[34] The probability of inserting at a uniformly chosen time and location is $d\tau/(V\beta)$ while the probability of removing one of $k+1$ vertices is $1/(k+1)$. The infinitesimal $d\tau$ cancels with infinitesimals in the weights, the a-priori probability $1/2$ of choosing one of the two values of the auxiliary field cancels with a factor $1/2$ in Eq. (10) and the factor $\beta V/(k+1)$ enters the acceptance ratio.